

A hybrid ensemble transform particle filter for nonlinear and spatially extended dynamical systems

Nawinda Chustagulprom*

Sebastian Reich†

Maria Reinhardt ‡

February 26, 2016

Abstract

Data assimilation is the task to combine evolution models and observational data in order to produce reliable predictions. In this paper, we focus on ensemble-based recursive data assimilation problems. Our main contribution is a hybrid filter that allows one to adaptively bridge between ensemble Kalman and particle filters. While ensemble Kalman filters are robust and applicable to strongly nonlinear systems even with small and moderate ensemble sizes, particle filters are asymptotically consistent in the large ensemble size limit. We demonstrate numerically that our hybrid approach can improve the performance of both Kalman and particle filters at moderate ensemble sizes. We also show how to implement the concept of localization into a hybrid filter, which is key to its applicability to spatially extended systems.

Keywords. Bayesian inference, Monte Carlo method, sequential data assimilation, ensemble Kalman filter, particle filter, localization, dynamical systems.

AMS(MOS) subject classifications. 65C05, 62M20, 93E11, 62F15, 86A22

1 Introduction

This paper is concerned with algorithms for assimilating observational data into computational models using ensemble prediction techniques and sequential Bayesian model updates. The focus will be on state estimation problems which give rise to the classic filtering problem of stochastic processes [14].

There are two popular classes of ensemble based filtering techniques; namely particle filters (PFs) [9] and ensemble Kalman filters (EnKFs) [10]. While PFs lead to consistent approximations of the underlying filtering problem [2], they suffer from the curse of dimensionality [4]. EnKFs, on the other hand, are applicable to spatially extended systems when combined with inflation and localization [10]; but they lead to estimates which are inconsistent with the Bayesian approach to filtering unless all involved probability density functions (PDFs) are Gaussian. There have been several recent attempts to enhance the performance of PFs by either using better proposal densities (see [31] for a review) or by implementing localisation into PFs [24, 6]. At the same time, several efforts have been made to extend EnKFs beyond their Gaussian prior assumption. Most of these attempts are built around Gaussian mixture approximations and, hence, can be seen as efforts to bridge EnKFs with PFs (see, for example, [28, 29, 25, 11]). Other recent attempts to extend EnKFs beyond Gaussianity include [1, 17, 21].

In this paper, we combine the general linear ensemble transform filter (LETF) framework, as laid out in detail in [27], with the EnKF-PF bridging approach suggested in [11]. As in [11], our hybrid ensemble filter approach bridges EnKFs with PFs through a parameter $\alpha \in [0, 1]$, which can be chosen adaptively. Key innovations of our approach, as compared to the hybrid method of [11], include the use of R -localisation [13, 27], the fact that either a PF or an EnKF can be applied first, that any implementation of an EnKF can be used (including [1]), and that observation operators need not to be linear.

The properties of our hybrid ensemble filter will be demonstrated numerically for the Lorenz-63 model [18], the Lorenz-96 model [19, 20], and a coupled Lorenz-96 wave equation model [5]. The numerical results

*Universität Potsdam, Institut für Mathematik, Karl-Liebknecht-Str. 24/25, D-14476 Potsdam, Germany

†Universität Potsdam, Institut für Mathematik, Karl-Liebknecht-Str. 24/25, D-14476 Potsdam, Germany (sreich@math.uni-potsdam.de) and University of Reading, Department of Mathematics and Statistics, Whiteknights, PO Box 220, Reading RG6 6AX, UK.

‡Universität Potsdam, Institut für Mathematik, Karl-Liebknecht-Str. 24/25, D-14476 Potsdam, Germany

demonstrate that our hybrid approach leads to reduced root mean square (RMS) errors compared to an EnKF even for ensemble sizes at which a standard PF fails to track the reference solution.

We now summarize the basic data assimilation problem which will be considered throughout this paper. See [16, 27] for further background material on filtering and general data assimilation algorithms. The basic task of sequential filtering (or sequential data assimilation) is to estimate a reference trajectory $z_{\text{ref}}(t) \in \mathbb{R}^{N_z}$ of an evolution equation, such as

$$\frac{dz}{dt} = f(z, t), \quad (1)$$

from partial and noisy observations

$$y_{\text{obs}}(t_k) = h(z_{\text{ref}}(t_k), t) + \xi_k, \quad k = 1, 2, \dots, \quad (2)$$

at discrete times $t_k > 0$. Here $h : \mathbb{R}^{N_z} \times \mathbb{R}_{>0} \rightarrow \mathbb{R}^{N_y}$ is the observation operator and $\xi_k \in \mathbb{R}^{N_y}$, $k \geq 1$, denote independent realizations of a Gaussian random variable with mean zero and covariance matrix $R \in \mathbb{R}^{N_y \times N_y}$. We are also given the PDF, $\pi_Z(z, 0)$, for the initial conditions $z(0)$ of (1) at time zero. In order to simplify the subsequent discussion, we will assume that the observation operator h is linear and time independent, i.e., $h(z, t) = Hz$.

EnKFs and PFs both aim at providing approximations to the conditional distributions

$$\pi_Z(z, t | y_{\text{obs}}(t_1), \dots, y_{\text{obs}}(t_k))$$

for $t \geq t_k$ in the form of an ensemble $z_i(t)$, $i = 1, \dots, M$. The best estimate for $z_{\text{ref}}(t)$ is provided by the ensemble mean,

$$z_{\text{est}}(t) = \frac{1}{M} \sum_{i=1}^M z_i(t). \quad (3)$$

While $z_{\text{est}}(t)$ from a PF converges to the true expectation value

$$\bar{z}(t) = \int_{\mathbb{R}^{N_z}} z \pi_Z(z, t | y_{\text{obs}}(t_1), \dots, y_{\text{obs}}(t_k)) dz \quad (4)$$

as $M \rightarrow \infty$, the same does not hold true for EnKFs, in general. However, it is often found in practice that EnKFs lead to smaller time-averaged RMS errors,

$$\text{RMSE} = \frac{1}{T} \int_0^T \sqrt{\frac{1}{N_z} \|z_{\text{est}}(t) - z_{\text{ref}}(t)\|^2} dt, \quad (5)$$

for small or moderate ensemble sizes, M . See [8] for more details.

Any ensemble-based filter algorithm alternates between prediction steps, where an ensemble is propagated under the dynamics (1), and analysis steps, where a forecast ensemble $z_i^f = z_i(t_k)$ is transformed into an analysis ensemble z_i^a based on the available observations at time $t = t_k$. The analysis ensemble then provides the initial conditions for the next prediction step over the time interval $t \in (t_k, t_{k+1}]$. It is the implementation of the analysis step that provides the crucial difference between the various filter algorithms. Section 2 lays out a general framework for formulating analysis steps and discusses the ensemble square root filter (ESRF) [30] and the ensemble transform particle filter (ETPF) [6] as two specific examples for EnKFs and PFs, respectively, which will be used throughout this paper. The novel hybrid filter is based on a combination of the ESRF and the ETPF and is introduced in detail in Section 2.3. Its generalization to spatially extended systems using localization is formulated in Section 3. The complete hybrid algorithm is summarized in Section 4, while the numerical results can be found in Section 5.

2 Summary of the LETF framework

Given a forecast ensemble z_i^f , $i = 1, \dots, M$, at an observation time t_k , we consider general linear transformations [27]

$$z_j^a = \sum_{i=1}^M z_i^f d_{ij} \quad (6)$$

into an analysis ensemble z_j^a , $j = 1, \dots, M$. The coefficients d_{ij} satisfy

$$\sum_{i=1}^M d_{ij} = 1 \quad (7)$$

and depend on the forecast ensemble and the observed $y_{\text{obs}} = y_{\text{obs}}(t_k)$. We demonstrate in the next two paragraphs how the ESRF and the ETPF fit into the LETF framework. See [27] for further details and LETF formulations of other EnKFs in particular.

2.1 ESRF formulation

We introduce the forecast ensemble mean

$$\bar{z}^f = \frac{1}{M} \sum_{i=1}^M z_i^f, \quad (8)$$

the $N_z \times M$ matrix of ensemble anomalies

$$A^f = [(z_1^f - \bar{z}^f) \ (z_2^f - \bar{z}^f) \ \dots \ (z_M^f - \bar{z}^f)], \quad (9)$$

the symmetric $M \times M$ matrix

$$S = \left\{ I + \frac{1}{M-1} (HA^f)^T R^{-1} HA^f \right\}^{-1/2}, \quad (10)$$

and the weight vector $\hat{w} \in \mathbb{R}^M$ with entries

$$\hat{w}_i = \frac{1}{M} - \frac{1}{M-1} e_i^T S^2 (HA^f)^T R^{-1} (H\bar{z}^f - y_{\text{obs}}). \quad (11)$$

Here e_i denotes the i th basis vector in \mathbb{R}^M . Also note that the entries s_{ij} of the matrix S satisfy

$$\sum_{j=1}^M s_{ij} = \sum_{i=1}^M s_{ij} = 1 \quad \text{and} \quad \sum_{i=1}^M \hat{w}_i = 1. \quad (12)$$

Following the presentation in [27], the analysis ensemble of an ESRF is now given by

$$z_j^a = \sum_{i=1}^M z_i^f \hat{w}_i + \sum_{i=1}^M (z_i^f - \bar{z}^f) s_{ij} = \sum_{i=1}^M z_i^f d_{ij}^{\text{KF}} \quad (13)$$

with filter coefficients

$$d_{ij}^{\text{KF}} = d_{ij}^{\text{KF}}(\{z_i^f\}, y_{\text{obs}}) := \hat{w}_i - \frac{1}{M} + s_{ij}, \quad (14)$$

which depend on the forecast ensemble and the observation, y_{obs} . We note that the analysis mean satisfies

$$\bar{z}^a = \sum_{i=1}^M z_i^f \hat{w}_i. \quad (15)$$

In case of nonlinear observation operators, one would replace HA^f by the $N_y \times M$ matrix

$$[(h(z_1^f, t_k) - \bar{h}^f) \ (h(z_2^f, t_k) - \bar{h}^f) \ \dots \ (h(z_M^f, t_k) - \bar{h}^f)] \quad (16)$$

and $H\bar{z}^f$ by

$$\bar{h}^f = \frac{1}{M} \sum_{i=1}^M h(z_i^f, t_k) \quad (17)$$

in the formulas above. We finally mention that the computational complexity of the ESRF scales like $\mathcal{O}(M^3)$ in the ensemble size M .

2.2 ETPF formulation

The ETPF provides an implementation of PFs which replaces the resampling step of a sequential importance resampling (SIR) PF [9] by a linear transformation of the form (6) [26, 6, 27]. As for a SIR PF, we define importance weights

$$w_i = \frac{\exp\left(-\frac{1}{2}(Hz_i^f - y_{\text{obs}})^T R^{-1}(Hz_i^f - y_{\text{obs}})\right)}{\sum_{j=1}^M \exp\left(-\frac{1}{2}(Hz_j^f - y_{\text{obs}})^T R^{-1}(Hz_j^f - y_{\text{obs}})\right)}. \quad (18)$$

The required transformation (6) is then found by solving the linear transport problem

$$T^* = \arg \min_{T \in \Pi} \sum_{i,j=1}^M t_{ij} \|z_i^f - z_j^f\|^2 \quad (19)$$

where Π denotes the set of $M \times M$ matrices T with non-negative entries t_{ij} subject to

$$\sum_{i=1}^M t_{ij} = 1, \quad \text{and} \quad \sum_{j=1}^M t_{ij} = w_i M. \quad (20)$$

Finally, the ETPF defines the analysis ensemble via

$$z_j^a = \sum_{i=1}^M z_i^f d_{ij}^{\text{PF}} \quad (21)$$

with filter coefficients

$$d_{ij}^{\text{PF}} = d_{ij}^{\text{PF}}(\{z_l^f\}, y_{\text{obs}}) := t_{ij}^* \quad (22)$$

and analysis ensemble mean

$$\bar{z}^a = \frac{1}{M} \sum_{j=1}^M z_j^a = \frac{1}{M} \sum_{j=1}^M \sum_{i=1}^M z_i^f t_{ij}^* = \sum_{i=1}^M z_i^f w_i. \quad (23)$$

The consistency of a single ETPF step with Bayes' formula has been demonstrated in [26]. However, the ETPF can fail to track the reference solution, $z_{\text{ref}}(t)$, if the effective sample size

$$M_{\text{eff}} = \frac{1}{\sum_{i=1}^M w_i^2} \quad (24)$$

becomes too small. In the extreme case of $M_{\text{eff}} = 1$, all analysis ensemble members, z_j^a , become equal to $z_{i^*}^f$, where $z_{i^*}^f$ is the forecast ensemble member with importance weight $w_{i^*} = 1$. This limitation of the ETPF provides the main motivation for our hybrid ensemble transform particle filter.

The linear transport problem (19) can be solved by a linear programming algorithm [22]. In our numerical experiments we used the *FastEMD* algorithm of [23], which has a computational complexity of $\mathcal{O}(M^3 \log M)$. We also note that one dimensional transport problems can be solved efficiently by sorting the ensemble members [27].

Nonlinear observation operators can be accommodated by the ETPF by simply replacing $H z_i^f$ by $h(z_i^f, t_k)$ in the computation of the importance weights (18).

2.3 The hybrid ensemble transform particle filter

Following [11], we split the likelihood function into a product of two factors, i.e.,

$$\pi_Y(y_{\text{obs}}|z) \propto \exp\left(-\frac{\alpha}{2}(Hz - y_{\text{obs}})^T R^{-1}(Hz - y_{\text{obs}})\right) \times \exp\left(-\frac{1-\alpha}{2}(Hz - y_{\text{obs}})^T R^{-1}(Hz - y_{\text{obs}})\right) \quad (25)$$

for $\alpha \in [0, 1]$. The second factor is treated by the ESRF giving rise to

$$S(\alpha) = \left\{ I + \frac{1-\alpha}{M-1} (HA^f)^T R^{-1} HA^f \right\}^{-1/2}, \quad (26)$$

weight vector $\hat{w}(\alpha) \in \mathbb{R}^M$ with entries

$$\hat{w}_i(\alpha) = \frac{1}{M} - \frac{1-\alpha}{M-1} e_i^T S(\alpha)^2 (H A^f)^T R^{-1} (H z^f - y_{\text{obs}}), \quad (27)$$

and filter coefficients

$$d_{ij}^{\text{KF}}(\alpha, \{z_l^f\}, y_{\text{obs}}) := \hat{w}_i(\alpha) - \frac{1}{M} + s_{ij}(\alpha). \quad (28)$$

The first factor is treated by the ETPF with importance weights

$$w_i(\alpha) = \frac{\exp\left(-\frac{\alpha}{2} (H z_i^f - y_{\text{obs}})^T R^{-1} (H z_i^f - y_{\text{obs}})\right)}{\sum_{j=1}^M \exp\left(-\frac{\alpha}{2} (H z_j^f - y_{\text{obs}})^T R^{-1} (H z_j^f - y_{\text{obs}})\right)} \quad (29)$$

and the linear transport problem

$$T^*(\alpha) = \arg \min_{T \in \Pi} \sum_{i,j=1}^M t_{ij} \|z_i^f - z_j^f\|^2, \quad (30)$$

where Π now denotes the set of $M \times M$ matrices T with $t_{ij} \geq 0$ subject to

$$\sum_{i=1}^M t_{ij} = 1, \quad \text{and} \quad \sum_{j=1}^M t_{ij} = w_i(\alpha) M. \quad (31)$$

The coefficients of the associated ETPF are then given by

$$d_{ij}^{\text{PF}}(\alpha, \{z_l^f\}, y_{\text{obs}}) := t_{ij}^*(\alpha). \quad (32)$$

When implementing the hybrid filter, one can decide between two options; either one applies the ETPF or the ESRF first. The two resulting implementations can be summarized as follows:

(A) ETPF-ESRF

$$z_j^h = \sum_{i=1}^M z_i^f d_{ij}^{\text{PF}}, \quad d_{ij}^{\text{PF}} := d_{ij}^{\text{PF}}(\alpha, \{z_l^f\}, y_{\text{obs}}), \quad (33)$$

$$z_j^a = \sum_{i=1}^M z_i^h d_{ij}^{\text{KF}}, \quad d_{ij}^{\text{KF}} := d_{ij}^{\text{KF}}(\alpha, \{z_l^h\}, y_{\text{obs}}). \quad (34)$$

(B) ESRF-ETPF

$$z_j^h = \sum_{i=1}^M z_i^f d_{ij}^{\text{KF}}, \quad d_{ij}^{\text{KF}} := d_{ij}^{\text{KF}}(\alpha, \{z_l^f\}, y_{\text{obs}}), \quad (35)$$

$$z_j^a = \sum_{i=1}^M z_i^h d_{ij}^{\text{PF}}, \quad d_{ij}^{\text{PF}} := d_{ij}^{\text{PF}}(\alpha, \{z_l^h\}, y_{\text{obs}}). \quad (36)$$

Note that (34) uses the updated ensemble from (33) to determined the coefficients d_{ij}^{KF} . The same dependency applies to the ESRF-ETPF.

The bridging parameter α can either be set to a constant value or selected adaptively [11] such that the ratio between the effective sample size,

$$M_{\text{eff}}(\alpha) := \frac{1}{\sum_{i=1}^M w_i(\alpha)^2}, \quad (37)$$

in the ETPF part of the hybrid filter and the ensemble size, M , is equal to a fixed reference value $\theta \in [0, 1]$, i.e.

$$\frac{M_{\text{ref}}(\alpha)}{M} = \theta. \quad (38)$$

Note that $\theta = 1$ implies $\alpha = 0$ and $M_{\text{ref}}(\alpha)$ is monotonically decreasing in α . If $M_{\text{ref}}(\alpha) > \theta M$ for all $\alpha \in [0, 1]$, then we set $\alpha = 1$.

2.4 Ensemble inflation and particle rejuvenation

Filter implementations with small or moderate ensemble sizes often lead to filter divergence, i.e., the filter is no longer able to track the reference solution. Filter divergence can often be traced back to analysis ensembles becoming underdispersive, i.e., the analysis ensemble spread does not properly represent Bayesian posterior uncertainties. In order to prevent ensembles from becoming underdispersive, one can apply multiplicative ensemble inflation to the forecast ensemble prior to an analysis step and/or one can apply particle rejuvenation to the analysis ensemble (see, for example, [10, 27]). We used particle rejuvenation in the form of

$$z_j^a \rightarrow z_j^a + \sum_{i=1}^M (z_i^f - \bar{z}^f) \frac{\beta \xi_{ij}}{\sqrt{M-1}} \quad (39)$$

in our numerical experiments. Here $\beta > 0$ denotes the rejuvenation parameter and the ξ_{ij} 's are i.i.d. Gaussian random variables with mean zero and variance one. We also enforce

$$\sum_{j=1}^M \xi_{ij} = 0 \quad (40)$$

in order to preserve the ensemble mean, \bar{z}^a , under rejuvenation.

Note that the particle rejuvenation step is the only instance in the complete filter algorithm where randomness enters. This randomness provides a safeguard against the creation of identically analysis ensemble members under an ETPF, as discussed towards the end of Section 2.2.

3 Spatially extended systems and localization

At each data assimilation step, we have to now deal with ensembles of spatial fields $z(x) \in \mathbb{R}^N$ where $x \in \mathbb{R}^d$ denotes the spatial variable of dimension d . Those fields are provided on a computational grid with grid points $x_k \in \mathbb{R}^d$, $k = 1, \dots, K$. Spatially discrete fields are formally collected into a state vector

$$z = (z(x_1)^T \ z(x_2)^T \ \dots \ z(x_K)^T)^T \quad (41)$$

of dimension $N_z = N \times K$. A forecast ensemble

$$z_i^f = (z_i^f(x_1)^T \ z_i^f(x_2)^T \ \dots \ z_i^f(x_K)^T)^T \in \mathbb{R}^{N_z}, \quad (42)$$

$i = 1, \dots, M$, is then transformed at each grid point, x_k , according to

$$z_j^a(x_k) = \sum_{i=1}^M z_i^f(x_k) d_{ij}(x_k), \quad d_{ij}(x_k) = d_{ij}(x_k, \{z_i^f\}, y_{\text{obs}}), \quad (43)$$

which gives rise to the analysis ensemble

$$z_j^a = (z_j^a(x_1)^T \ z_j^a(x_2)^T \ \dots \ z_j^a(x_K)^T)^T \in \mathbb{R}^{N_z}, \quad (44)$$

$j = 1, \dots, M$. It is crucial that these local updates are implemented in a manner that maintains the spatial regularity of the forecast states. This prevents the use of localization in a standard SIR PF.

We assume that observations are taken at a discrete subset of grid points \mathbf{x}_q , $q = 1, \dots, Q$, and satisfy the observation model

$$y_q = \mathcal{H}z(\mathbf{x}_q) + \eta_q$$

with independent measurement errors $\eta_q \sim \mathcal{N}(0, r)$ and observation operator $\mathcal{H} : \mathbb{R}^N \rightarrow \mathbb{R}^{N_y}$. We assume here, for simplicity, $N_y = 1$ and $r > 0$ denotes the variance of the measurement error. We introduce the more compact notation

$$y_{\text{obs}} = Hz + \eta \quad (45)$$

with $y_{\text{obs}} = (y_1, \dots, y_Q)^T \in \mathbb{R}^Q$, $\eta = (\eta_1, \dots, \eta_Q)^T \sim \mathcal{N}(0, R)$, $R = rI$, and $H : \mathbb{R}^{N_z} \rightarrow \mathbb{R}^Q$ appropriately defined. If each grid point is observed then we simply have $Q = K$, $\mathbf{x}_q = x_q$ with $q = 1, \dots, K$. If every second grid point is observed then $Q = K/2$ and $\mathbf{x}_q = x_{2q}$ with $q = 1, \dots, K/2$ etc.

We now extend the concept of R -localization [13, 6, 27] to the hybrid filters introduced in Section 2.3. First we define a localized inverse measurement error covariance matrix $R(x_k)^{-1} \in \mathbb{R}^{Q \times Q}$ with diagonal entries

$$r_{qq}(x_k)^{-1} = \frac{\rho\left(\frac{\|x_k - \mathbf{x}_q\|}{R_{\text{loc}}}\right)}{r}, \quad (46)$$

$q = 1, \dots, Q$, for given localization radius $R_{\text{loc}} > 0$ and compactly supported tempering function $\rho(t)$ with $\rho(0) = 1$ and $\rho(t) = 0$ for $t \geq 2$. We used

$$\rho(t) = \begin{cases} 1 - \frac{5}{3}t^2 + \frac{5}{8}t^3 + \frac{1}{2}t^4 - \frac{1}{4}t^5 & \text{for } t \leq 1, \\ -\frac{2}{3}t^{-1} + 4 - 5t + \frac{5}{3}t^2 + \frac{5}{8}t^3 - \frac{1}{2}t^4 + \frac{1}{12}t^5 & \text{for } 1 \leq t \leq 2, \\ 0 & \text{otherwise,} \end{cases} \quad (47)$$

in our numerical experiments [12]. The two hybrid algorithms from Section 2.3 are now adjusted as follows. The ESRF is replaced by the LETKF [13], which leads to

$$S(x_k, \alpha_k) = \left\{ I + \frac{1 - \alpha_k}{M - 1} (HA^f)^T R(x_k)^{-1} HA^f \right\}^{-1/2}, \quad (48)$$

$$\hat{w}_i(x_k, \alpha_k) = \frac{1}{M} - \frac{1 - \alpha_k}{M - 1} e_i^T S(x_k, \alpha_k)^2 (HA^f)^T R(x_k)^{-1} (Hz^f - y_{\text{obs}}), \quad (49)$$

and

$$d_{ij}^{\text{KF}}(x_k, \alpha_k, \{z_i^f\}, y_{\text{obs}}) := \hat{w}_i(x_k, \alpha_k) - \frac{1}{M} + s_{ij}(x_k, \alpha_k). \quad (50)$$

The ETPF filter contribution is modified at each grid point x_k as proposed in [6, 27]. Specifically, we define localized weights

$$w_i(x_k, \alpha_k) = \frac{\exp\left(-\frac{\alpha_k}{2} (Hz_i^f - y_{\text{obs}})^T R(x_k)^{-1} (Hz_i^f - y_{\text{obs}})\right)}{\sum_{j=1}^M \exp\left(-\frac{\alpha_k}{2} (Hz_j^f - y_{\text{obs}})^T R(x_k)^{-1} (Hz_j^f - y_{\text{obs}})\right)}, \quad (51)$$

introduce the linear transport problem

$$T^* = \arg \min_{T \in \Pi} \sum_{i,j=1}^M t_{ij} \|z_i^f(x_k) - z_j^f(x_k)\|^2 \quad (52)$$

at each grid point, x_k , where Π denotes the set of $M \times M$ matrices T with $t_{ij} \geq 0$ subject to

$$\sum_{i=1}^M t_{ij} = 1, \quad \text{and} \quad \sum_{j=1}^M t_{ij} = w_i(x_k, \alpha_k) M, \quad (53)$$

and define

$$d_{ij}^{\text{PF}}(x_k, \alpha_k, \{z_i^f\}, y_{\text{obs}}) = t_{ij}^*. \quad (54)$$

We note that the distance term in (52) can be replaced by other distance measures involving, for example, grid point values in the vicinity of x_k . See [6, 27] for more details. The cost functional (52) has the advantage that the associated transport problem is easy to solve in case $z(x_k) \in \mathbb{R}$, i.e. $N = 1$.

Finally, the desired localized hybrid filter schemes can be implemented in the following two variants:

(A) ETPF-LETKF

$$z_j^{\text{h}}(x_k) = \sum_{i=1}^M z_i^f(x_k) d_{ij}^{\text{PF}}(x_k), \quad d_{ij}^{\text{PF}}(x_k) := d_{ij}^{\text{PF}}(x_k, \alpha_k, \{z_i^f\}, y_{\text{obs}}), \quad \text{over all grid points,} \quad (55)$$

$$z_j^{\text{a}}(x_k) = \sum_{i=1}^M z_i^{\text{h}}(x_k) d_{ij}^{\text{KF}}(x_k), \quad d_{ij}^{\text{KF}}(x_k) := d_{ij}^{\text{KF}}(x_k, \alpha_k, \{z_i^{\text{h}}\}, y_{\text{obs}}), \quad \text{over all grid points.} \quad (56)$$

(B) LETKF-ETPF

$$z_j^h(x_k) = \sum_{i=1}^M z_i^f(x_k) d_{ij}^{\text{KF}}(x_k), \quad d_{ij}^{\text{KF}}(x_k) := d_{ij}^{\text{KF}}(x_k, \alpha_k, \{z_l^f\}, y_{\text{obs}}), \quad \text{over all grid points}, \quad (57)$$

$$z_j^a(x_k) = \sum_{i=1}^M z_i^h(x_k) d_{ij}^{\text{PF}}(x_k), \quad d_{ij}^{\text{PF}}(x_k) := d_{ij}^{\text{PF}}(x_k, \alpha_k, \{z_l^h\}, y_{\text{obs}}), \quad \text{over all grid points}. \quad (58)$$

The bridging parameters, α_k , can again be either set to a constant value, α , for all grid points or determined adaptively. In the latter case α_k is determined by

$$\frac{M_{\text{ref}}(x_k \alpha_k)}{M} = \theta, \quad M_{\text{ref}}(x_k, \alpha_k) := \frac{1}{\sum_{i=1}^M w_i(x_k, \alpha_k)^2}, \quad (59)$$

for given effective ensemble size ratio, $\theta \in [0, 1]$. If $M_{\text{ref}}(x_k, \alpha_k) > \theta M$ for all $\alpha_k \in [0, 1]$, then we set $\alpha_k = 1$. Finally, particle rejuvenation in the form of (39) is performed at the end of each filter step.

4 Summary of hybrid filter algorithm

We summarize the numerical implementation of the proposed hybrid filter as it has been used for the numerical experiments below. We explicitly describe the filter implementation with localization. An implementation without localization follows from straightforward modifications. The free parameters of the data assimilation steps are the ensemble size, $M > 1$, the localization radius, $R_{\text{loc}} > 0$, in (46), the bridging parameters, $\alpha_k \in [0, 1]$, which are either fixed to a value α or determined adaptively through the effective sample size ratio, $\theta \in [0, 1]$, in (59). Any choice of these parameters together with a forecast ensemble z_l^f , $l = 1, \dots, M$, and an observation $y_{\text{obs}}(t_k)$ determine the ETPF-LETKF update step (55)-(56) uniquely. The linear transport problems, defined by (52)-(53) are solved using the MATLAB implementation of *FastEMD* [23].

The update step (55)-(56) is followed by particle rejuvenation (39) with parameter $\beta > 0$. Note that rejuvenation is applied to the whole spatially dependent analysis fields with one and the same set of random numbers ξ_{ij} , i.e.

$$z_j^a(x_k) \rightarrow z_j^a(x_k) + \sum_{i=1}^M (z_i^f(x_k) - \bar{z}^f(x_k)) \frac{\beta \xi_{ij}}{\sqrt{M-1}} \quad (60)$$

for all grid points x_k .

A complete implementation also requires a time-stepping method for the differential equation (1). We used the implicit midpoint rule with step-size denoted by $\Delta t > 0$. Observations (2) are generated from a reference trajectory, $\{z_{\text{ref}}^n\}_{n \geq 0}$, which is computed using the same time-stepping method in (1) and a randomly selected initial condition $z_{\text{ref}}^0 \sim \pi_Z(z, 0)$. The initial ensemble, z_i^0 , $i = 1, \dots, M$, is also drawn from the initial PDF, $\pi_Z(z, 0)$.

5 Numerical experiments

We now conduct a series of numerical experiments in order to assess the relative merits of the hybrid approach compared to a purely EnKF or PF based data assimilation approach. The ensemble sizes, M , are chosen such that standard PF implementations cannot compete with an EnKF. Such a data assimilation scenario is typical in complex application areas such as numerical weather prediction [15]. The quality of a filter is assessed by the following time-discretised version of (5):

$$\text{RMSE} = \frac{1}{K} \sum_{i=1}^K \sqrt{\frac{1}{N_z} \|\bar{z}^a(t_k) - z_{\text{ref}}(t_k)\|^2}. \quad (61)$$

Here K denotes the total number of assimilation steps.




Figure 1: Hybrid filter for Lorenz-63 system with ensemble sizes $M = 15, \dots, 35$: left panel: RMS errors for fixed bridging parameter; right panel: RMS errors for adaptively determined bridging parameter. Note that $\alpha = 0$ and $\theta = 1$ correspond to a pure ESRF filter.

5.1 Lorenz-63

The first numerical tests of our hybrid filter are performed on the chaotic Lorenz-63 system [18] with its standard parameter setting $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$. We only observe the first component of the three dimensional system in observation intervals of $\Delta t_{\text{obs}} = 0.12$, i.e., $t_k = 0.12k$ in (61), and with observation error variance $R = 8$. A total of $K = 100,000$ assimilation steps are performed and $N_z = 3$ in (61). The differential equations are solved numerically with a step-size of $\Delta t = 0.01$.

The ETPF-ESRF filter is implemented for ensemble sizes varying between $M = 15$ and $M = 35$. Particle rejuvenation is applied with $\beta = 0.2$. Simulations with $\beta = 0.15$ and $\beta = 0.25$ gave similar results to those reported here. No localization is necessary for this low dimensional system. Our numerical results for fixed and adaptively determined bridging parameter, α , can be found in Figure 1. It can be seen that a fixed bridging parameter of $\alpha = 0.2$ leads already to substantial improvements for the smallest ensemble size of $M = 15$. Further improvements are achieved for larger ensemble sizes at larger bridging parameters. In contrast to these very encouraging results for fixed bridging parameters, the adaptive implementation of the hybrid filter does not work quite as effectively for small ensemble sizes of $M = 15$ and $M = 20$. This finding indicates that one might have to search for other selection criteria.

We also implemented the Ensemble Kalman Particle Filter of [11] for comparison. The Ensemble Kalman Particle Filter bridges the EnKF with perturbed observations [10] and a SIR PF. Here we replaced the SIR PF with the ETPF in order to facilitate a fair comparison. The numerical results are displayed in Figure 2. We find that the Ensemble Kalman Particle Filter leads to larger RMS errors for all values of the bridging parameter, α . We speculate that this is due to the use of the EnKF with perturbed observations and the fact that the EnKF is applied first. This hypothesis is supported by the numerical results from the ESRF-ETPF, which also does not lead to significant improvements for $\alpha > 0$. See Figure 2.

Note that we run the hybrid filters at ensemble sizes, M , for which EnKFs outperform PFs. The more interesting it is that the ETPF-ESRF leads to substantial improvements for values of the bridging parameter in the range $[0.2, 0.4]$ for ensemble sizes between $M = 15$ and $M = 35$. Figure 1 shows that the optimal α increases with increasing ensemble sizes, M . Asymptotically, one expects that $\alpha \rightarrow 1$ and $\theta \rightarrow 0$ as $M \rightarrow \infty$. Indeed, we found numerically that $\alpha = 0$ is optimal for ensemble sizes smaller than $M = 10$ and that $\alpha = 1$ for ensemble sizes larger than $M = 100$. See also the convergence study in Section 5.4.

5.2 Lorenz-96

We next test our hybrid approach on the spatially extended Lorenz-96 model

$$\dot{x}_l = (x_{l+1} - x_{l-2})x_{l-1} - x_l + 8, \quad l = 1, \dots, 40 \quad (62)$$




Figure 2: Left panel: RMS errors for Ensemble Kalman Particle Filter applied to the Lorenz-63 system with ensemble sizes $M = 15, \dots, 35$ and for fixed bridging parameter; right panel: RMS errors for the ESRF-ETPF using the same setting and α in the range $[0, 0.5]$.

[19, 20]. We observe every other grid point in time intervals of $\Delta t_{\text{obs}} = 0.11$ with independent Gaussian measurement errors of variance equal to eight. A total of $K = 50,000$ assimilation steps is performed. The hybrid ETPF-LETKF filter is implemented with localization radius $R_{\text{loc}} = 4$, particle rejuvenation $\beta = 0.2$ and ensemble sizes $M = 20, 25, 30$. The equations of motion are integrated numerically with the implicit midpoint rule and $\Delta t = \Delta t_{\text{obs}}/22$. The resulting time-averaged RMSE can be found in Figure 3. The overall findings are similar to the Lorenz-63 model except that the adaptive approach now performs comparably (or even better) to an optimally tuned fixed parameter implementation and that the overall improvements in terms of RMS errors are less significant. The later fact is likely due to the weaker nonlinearity of the Lorenz-96 model. We also implemented the hybrid filter for localization radii $R_{\text{loc}} \in \{3, 5, 6, 7, 8\}$ and found that $\alpha > 0$ always leads to improvements in the RMS error compared to a pure LETKF implementation under the simulation conditions stated above. We finally implemented the Lorenz-96 in the more traditional setting of each grid point being observed in intervals of $\Delta t_{\text{obs}} = 0.05$ with measurement error variance $r = 1$. Our numerical findings indicate that the achievable improvements of our hybrid filter in comparison to the LETKF are rather minor and require ensemble sizes $M \geq 30$. We believe that these findings reflect the fact that frequent observations of the complete state space lead to nearly Gaussian forecast distributions.

5.3 Lorenz-96 coupled to a wave equation

Physical processes in the geosciences often act on different time-scales. To simulate that effect we use a model, that consists of the advective Lorenz-96 system (62) and relatively fast moving waves, which are modeled by the discrete wave equation

$$\epsilon^2 \ddot{h}_l = -h_l + c^2(h_{l+1} - 2h_l + h_{l-1}) - 2\epsilon^2 \gamma \dot{h}_l, \quad l = 1, \dots, 40. \quad (63)$$

The positive parameter $\epsilon \ll 1$ characterises the ratio between the advective velocity in x and the wave speed of h . The parameter $c > 0$ determines the wave dispersion over the computational grid and $\gamma > 0$ is a damping parameter. The coupled model with coupling-strength δ looks as follows [5]:

$$\dot{x}_l = (1 - \delta)(x_{l+1} - x_{l-2})x_{l-1} + \delta(x_{l-1}h_{l+1} - x_{l-2}h_{l-1}) - x_l + 8, \quad (64)$$

$$\epsilon^2 \ddot{h}_l = -h_l + c^2(h_{l+1} - 2h_l + h_{l-1}) + x_l - 2\epsilon^2 \gamma \dot{h}_l, \quad (65)$$

$l = 1, \dots, 40$, and periodic boundary conditions are applied. There exists a balance relation

$$x_l = h_l - c^2(h_{l+1} - 2h_l + h_{l-1}) \quad (66)$$

between the x and h variables. If it is fulfilled initially, solutions stay in approximate balance for long time intervals with the deviations from exact balance proportional to ϵ . For our experiments, we used this model

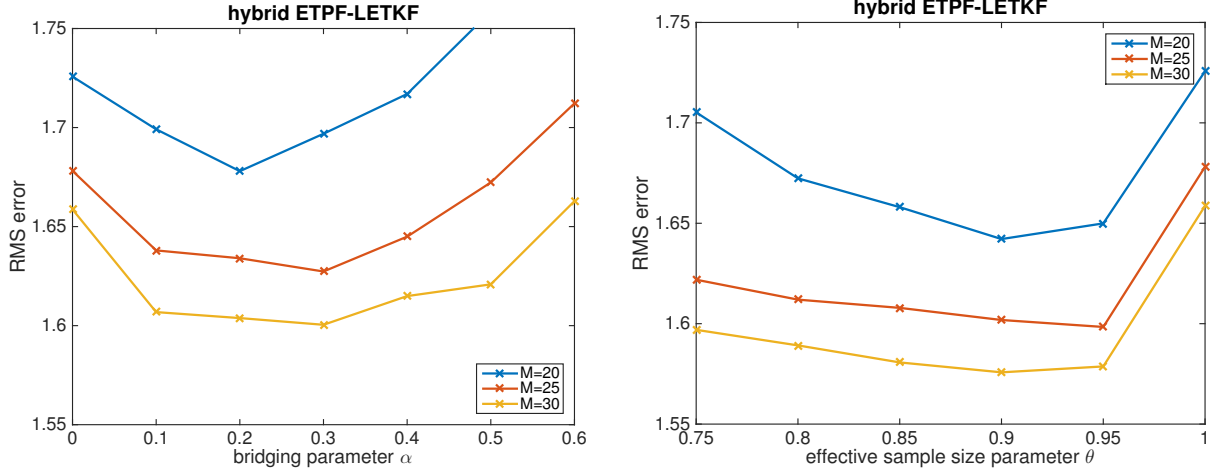


Figure 3: Hybrid filter for Lorenz-96 system with ensemble sizes $M = 20, 25, 30$: left panel: RMS errors for fixed bridging parameter; right panel: RMS errors for adaptively determined bridging parameters

with a coupling strength of $\delta = 0.1$, $\epsilon = 0.0025$, a damping of $\gamma = 0.1$, and $c = 1/2$. The time-stepping method of [5] is used with a step-size of $\Delta t = 0.002$.

The initial ensemble in $\{x_l\}$ is created by sampling from a Gaussian distribution with mean zero and covariance $0.01I$. The h -values of the initial ensemble are created by calculating them from the x -values using the balance relation (66). The \dot{h} -values are set equal to zero. We observe x_l at every other grid-point at time intervals of $\Delta t_{obs} = 0.15$ with measurement-error variance $R = 8I_{20}$. We applied the hybrid ETPF-LETKF filter to our model with a localization radius of $R_{loc} = 4$ and a particle-rejuvenation parameter of $\beta = 0.2$. We let the experiment run for $K = 50,000$ assimilation cycles and a spin-up phase of 1000 cycles, in which α is set to zero and the h -values are calculated to be in exact balance to the x -values after every assimilation step. We conduct experiments for ensemble sizes of $M = 20, 25, 30$, fixed bridging parameters $\alpha = 0, 0.1, \dots, 0.6$ and also for adaptive α_k 's using $\theta = 0.75, 0.8, \dots, 1$.

The resulting time-averaged RMS errors are displayed in Figure 4. One can see that the minimal RMS error is obtained at $\alpha = 0.3$ when $M = 20$, at $\alpha = 0.4$ when $M = 25$ and at $\alpha = 0.5$ when $M = 30$. The results for the adaptive approach show, that the RMS error is minimal for values of $\theta \in [0.90, 0.95]$ depending on the ensemble size, M , and that the adaptive approach is less efficient than an optimal choice of α . This is likely due to the fact that an adaptive choice of the bridging parameter in each grid point, x_k , together with localization leads to less balanced fields.

We finally mention that the damping in (65) is necessary to suppress the accumulation of imbalances generated in each assimilation step. See [5] for a detailed discussion of this phenomenon and a potential cure in the form of a mollified filter approach.

5.4 Convergence study for a single Bayesian inference step

After conducting a series of numerical experiments for which the ensemble sizes were relatively small, we finally demonstrate that the proposed hybrid filter indeed bridges EnKFs and PFs as the ensemble size varies from very small to very large. In order to keep the computations feasible, we only perform a single assimilation step with given bimodal prior (see Figure 5) and the ETPF-ESRF implemented for ensemble sizes $M \in [2, 4, 8, \dots, 1024]$. Each experiment is repeated 100,000 times and the averaged RMS errors for the estimated posterior mean are computed in terms of M and α . The optimal value of α and its RMS error are displayed in Figure 5 as a function of the ensemble size, M . As expected, the optimal value of α is close to zero for small M and approaches one as M increases.

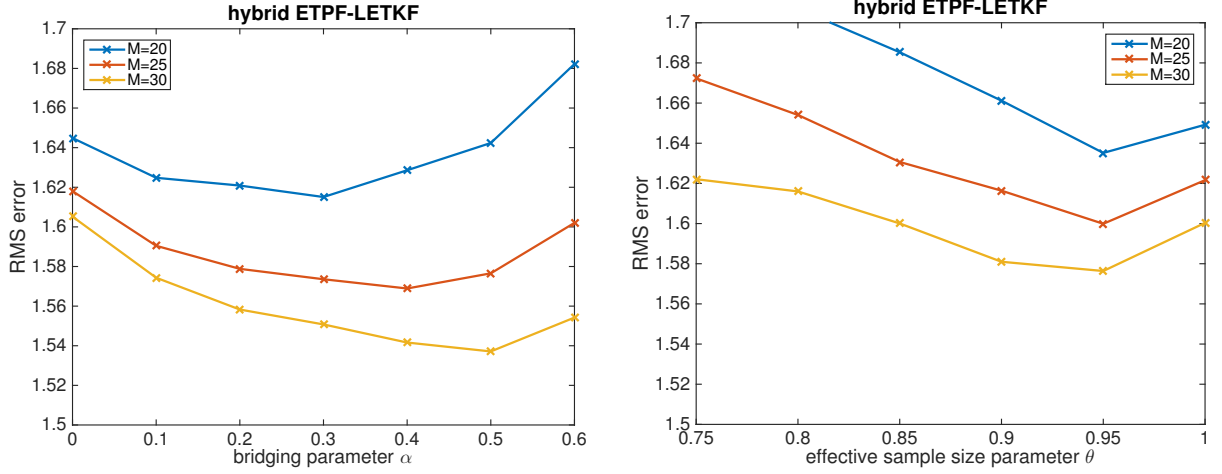


Figure 4: Hybrid filter for Lorenz-96 coupled to a wave equation with ensemble sizes $M = 20, 25, 30$: left panel: RMS errors for fixed bridging parameter; right panel: RMS errors for adaptively determined bridging parameters

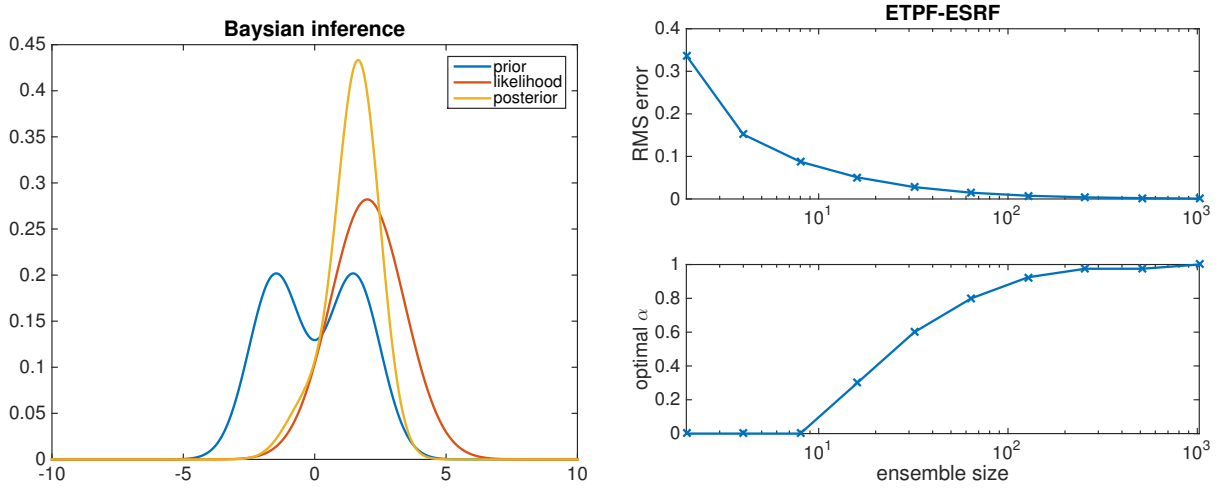


Figure 5: Single data assimilation step with ETPF-ESRF; left panel: prior and posterior distributions; right panel: optimal bridging parameter and resulting RMS errors as a function of the ensemble size.

6 Conclusions

Further numerical experiments show that the ETPF-ESRF and ETPF-LETKF, respectively, outperform the ESRF-ETPF and LETKF-ETPF, respectively, on all numerical examples of Section 5. See [7] for more details. We believe that this is due to the fact that the true forecast/prior distributions in our examples are more non-Gaussian than the analysis/posterior distributions. This is the case, in particular, for the Lorenz 63 model. Under those circumstances it seems preferable to apply a PF first followed by an EnKF. We also found that the adaptive hybrid method works well for the Lorenz-96 system, while a fixed choice of α is preferable for the Lorenz-63 model. It is feasible that other adaptation criteria for the choice of α could lead to better results. For example, one could select α such that the Kullback-Leibler divergence between the uniform measure, $1/M$, and the importance weights (29) remains below a given threshold.

Our numerical results also indicate that our hybrid method is suitable for multi-scale systems. However, our hybrid method does not address the problem of generating unbalanced fields through localization (see, for example, [3]). This issue will be addressed in a separate publication using the idea of mollification, as proposed in [5]. We finally mention that our hybrid methods can be combined with modified proposal densities for PFs (see [31] for a recent review of the subject).

Acknowledgments

SR and MR acknowledge support under the DFG Collaborative Research Center SFB 1114: Scaling Cascades in Complex Systems. NC's work has been supported by a fellowship of the Helmholtz Graduate School GeoSim. We thank Walter Acevedo for discussions on the implementation of the hybrid method.

References

- [1] J.L. Anderson. A non-Gaussian ensemble filter update for data assimilation. *Monthly Weather Review*, 138:4186–4198, 2010.
- [2] A. Bain and D. Crisan. *Fundamentals of stochastic filtering*, volume 60 of *Stochastic modelling and applied probability*. Springer-Verlag, New-York, 2008.
- [3] J. Ballabrera-Poy, E. Kalnay, and S.-C. Yang. Data assimilation in a system with two scales—combining two initialization techniques. *Tellus*, 61A:539–549, 2009.
- [4] T. Bengtsson, P. Bickel, and B. Li. Curse of dimensionality revisited: Collapse of the particle filter in very large scale systems. In *IMS Lecture Notes - Monograph Series in Probability and Statistics: Essays in Honor of David F. Freedman*, volume 2, pages 316–334. Institute of Mathematical Sciences, 2008.
- [5] K. Bergemann and S. Reich. A mollified ensemble Kalman filter. *Q. J. R. Meteorological Soc.*, 136:1636–1643, 2010.
- [6] Y. Chen and S. Reich. Assimilating data into scientific models: An optimal coupling perspective. In *Frontiers in Applied Dynamical Systems: Reviews and Tutorials*, volume 2, pages 75–118. Springer-Verlag, New York, 2015.
- [7] Nawinda Chutstagulprom. Phd thesis. Technical report, Universität Potsdam, Institut für Mathematik, Am Neuen Palais 10, 14469 Potsdam, 2016.
- [8] D. Crisan and B. Rozovskii, editors. *The Oxford Handbook of Nonlinear Filtering*. Oxford University Press, Oxford, 2011.
- [9] A. Doucet, N. de Freitas, and N. Gordon (eds.). *Sequential Monte Carlo methods in practice*. Springer-Verlag, Berlin Heidelberg New York, 2001.
- [10] G. Evensen. *Data assimilation. The ensemble Kalman filter*. Springer-Verlag, New York, 2006.
- [11] M. Frei and H.R. Künsch. Bridging the ensemble Kalman and particle filters. *Biometrika*, 100:781–800, 2013.

- [12] G. Gaspari and S.E. Cohn. Construction of correlation functions in two and three dimensions. *Q. J. Royal Meteorological Soc.*, 125:723–757, 1999.
- [13] B.R. Hunt, E.J. Kostelich, and I. Szunyogh. Efficient data assimilation for spatialtemporal chaos: A local ensemble transform Kalman filter. *Physica D*, 230:112–137, 2007.
- [14] A.H. Jazwinski. *Stochastic processes and filtering theory*. Academic Press, New York, 1970.
- [15] E. Kalnay. *Atmospheric modeling, data assimilation and predictability*. Cambridge University Press, 2002.
- [16] K. Law, A. Stuart, and K. Zygalakis. *Data Assimilation*. Springer-Verlag, New York, 2015.
- [17] J. Lei and P. Bickel. A moment matching ensemble filter for nonlinear and non-Gaussian data assimilation. *Mon. Weath. Rev.*, 139:3964–3973, 2011.
- [18] E.N. Lorenz. Deterministic non-periodic flows. *J. Atmos. Sci.*, 20:130–141, 1963.
- [19] E.N. Lorenz. Predictability: A problem partly solved. In *Proc. Seminar on Predictability*, volume 1, pages 1–18, ECMWF, Reading, Berkshire, UK, 1996.
- [20] E.N. Lorenz and K.E. Emanuel. Optimal sites for supplementary weather observations: Simulations with a small model. *J. Atmos. Sci.*, 55:399–414, 1998.
- [21] S. Metref, E. Cosme, C. Snyder, and P. Brasseur. A non-Gaussian analysis scheme using rank histograms for ensemble data assimilation. *Nonlinear Processes in Geophysics*, 21:869–885, 2013.
- [22] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer-Verlag, New York, 2nd edition, 2006.
- [23] O. Pele and M. Werman. Fast and robust earth mover’s distances. In *Computer Vision, 2009 IEEE 12th international conference*, pages 460–467, 2009.
- [24] P. Rebeschini and R. van Handel. Can local particle filters beat the curse of dimensionality? *Ann. Appl. Probab.*, 25:2809–2866, 2015.
- [25] S. Reich. A Gaussian mixture ensemble transform filter. *Q. J. R. Meteorolog. Soc.*, 138:222–233, 2012.
- [26] S. Reich. A nonparametric ensemble transform method for Bayesian inference. *SIAM J. Sci. Comput.*, 35:A2013–A2024, 2013.
- [27] S. Reich and C.J. Cotter. *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge University Press, Cambridge, 2015.
- [28] K.W. Smith. Cluster ensemble Kalman filter. *Tellus*, 59A:749–757, 2007.
- [29] A.S. Stordal, H.A. Karlsen, G. Nævdal, H.J. Skaug, and B. Vallés. Bridging the ensemble Kalman filter and particle filters: the adaptive Gaussian mixture filter. *Comput. Geosci.*, 15:293–305, 2011.
- [30] M.K. Tippett, J.L. Anderson, G.H. Bishop, T.M. Hamill, and J.S. Whitaker. Ensemble square root filters. *Mon. Wea. Rev.*, 131:1485–1490, 2003.
- [31] P.J. Van Leeuwen. Nonlinear data assimilation for high-dimensional systems. In *Frontiers in Applied Dynamical Systems: Reviews and Tutorials*, volume 2, pages 1–73. Springer-Verlag, New York, 2015.